IN-61
2-3023
15P

NASA Technical Memorandum 109058

A Programming Environment for Distributed Complex Computing

An Overview of the Framework for Interdisciplinary Design Optimization (FIDO) Project

NASA Langley TOPS Exhibit H120b

James C. Townsend and Robert P. Weston
Langley Research Center
Hampton, Virginia

Thomas M. Eidson
High Technology Corporation
Hampton, Virginia

DECEMBER 1993

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-0001

N94-24603
Unclas
G3/61 0205023
(NASA-TM-109058) A PROGRAMMING ENVIRONMENT FOR DISTRIBUTED COMPLEX COMPUTING. AN OVERVIEW OF THE FRAMEWORK FOR INTERDISCIPLINARY DESIGN OPTIMIZATION (FIDO) PROJECT. NASA LANGLEY TOPS EXHIBIT H120B (NASA) 15 p

ABSTRACT

The Framework for Interdisciplinary Design Optimization (FIDO) is a general programming environment for automating the distribution of complex computing tasks over a networked system of heterogeneous computers.

Many design problems require inputs from a number of specialty disciplines. (For airplanes, these include aerodynamics, structures, controls, etc.) The point of view, design emphasis, and design approach of each discipline specialist can be quite different. Designers frequently pass a new design from one discipline to another manually as each designer tries to reach an optimum with little direct interaction with the others. The FIDO system facilitates communications and provides for automatic interactions among the tasks in a multicomponent computational problem, such as a multidisciplinary design process, on a distributed heterogeneous computing system.

All computers involved are networked together, have access to centralized data, and work on their parts of the design simultaneously in parallel whenever possible. Each computational task is done by the computer type most appropriate for it. Provision is made for viewing results as they are produced and for steering the design process. The software is written in modular form to ease migration to upgraded or completely new problems; different codes can be substituted for each of the current code modules with little or no effect on the others.

The potential for commercial use of FIDO rests in its capability to automatically coordinate diverse computations on a networked system of workstations and computers. For example, FIDO could provide the coordination required for the design of fixed- or rotary-wing aircraft, automobiles, ships, spacecraft, computer systems, or electronics.

OVERVIEW OF THE FIDO SYSTEM

Introduction

Established as a part of the Langley High Performance Computation and Communication Program (HPCCP), the Framework for Interdisciplinary Design Optimization (FIDO) project has as its goal the development of a general programming environment for distributing a multicomponent computational problem across a networked system of heterogeneous computers. A multidisciplinary airplane design process was chosen for the development of the programming environment because of the interest in improving the efficiency of that process.

The FIDO system provides a means for automating the total design process. It facilitates communication and control between components of the system, which include the diverse discipline computations involved in a design problem and the system services that facilitate the design. The computers used can include workstations, vector supercomputers, and parallel-processing computers (figure 1), although only UNIX®-based workstations are used currently. Each computational task can be done by the computer type most appropriate for it. All of the computers involved are networked together, have access to centralized data, and work on their parts of the design, simultaneously whenever possible, under the coordination of a master code.

The following sections describe the model problem, the conceptual environment, task control, data management, discipline segment functionality, the communications library, and the results interrogation segment. The appendix gives a concise summary of the current project status.

FIDO Model Problem

A simple model of a High-Speed Civil Transport (HSCT) design problem (figure 2) was chosen for the initial implementation of the FIDO system. This problem is illustrative of the type of computational problem envisioned for the FIDO system; however, the scope is much reduced so that the focus can be specifically on computational system issues rather than on the design problem. Figure 3 is a view of the overall optimization loop for a simplified design of an HSCT. A particular design cycle begins with the choice of the flight conditions, base aircraft geometry, and the values of the design constraints and the design variables, as shown at the top of the figure.

For the FIDO model problem, four disciplines are involved in the analysis of the aircraft: aerodynamics, structures, performance, and propulsion. At each cycle of the design, these disciplines are invoked to analyze the current definition of the aircraft as specified by the current values of the design variables. The results of the analyses provide a set of system responses that correspond to the current design variables. The optimizer program uses these responses plus the derivatives of the system responses with respect to the design variables. (At the current stage of FIDO's development, these derivatives are obtained by finite differences using multiple

analyses on a perturbed set of design variables. In the future, they may be obtained by other methods, such as automatic differentiation within the discipline codes.)

In order to derive new values for the design variables, the optimizer uses an objective function (which in this case is the minimization of the aircraft's gross weight for the specified range and payload) along with the current system responses and derivatives. These new values are fed back into the design cycle until the process converges.

FIDO Environment

The conceptual environment in which the various discipline and service codes run is illustrated in figure 4. Each of the discipline codes (AERO, STRU, PERF, PROP) runs on a separate workstation, and they communicate with each other by sending and receiving information through a central data-manager code (DATA), which may run on a separate workstation. Because the role of the data manager is only to move information and manage its storage, a separate segment called the Interdisciplinary code (INTR) performs any computations necessary to convert data from one discipline into a form suitable for another discipline. (An example is the integration of aerodynamic pressures to provide structural forces.) Likewise, this code sends and receives its data from the data manager. A master program (MAST) controls the order in which the various discipline codes run. It also initializes the FIDO communication network (COMM Network), which forms the backbone of the system.

The discipline, data-manager, and master codes, along with the optimizer code (OPTI), comprise the primary computational segments of the FIDO network. To provide a means of looking at the data as the design process proceeds, an auxiliary segment called 'SPY' is added. The SPY segment can be started while computations are underway, and multiple instances of SPY are allowed. With proper permissions, these can run on any workstation connected to the network, which allows remote consultants to view results as they are produced and to give timely advice.

Not shown is an independent program (SETUP), which is invoked by the designer to pick the flight conditions, the base aircraft geometry, the initial design constraint values, and the initial values of the design variables from a range of previously stored possibilities.

Task Control

The master segment (MAST in figure 4) is responsible for the overall control flow of the primary segments in the execution system. A major exception is when multiple tasks (and, thus, segments) are started by the master in parallel; any synchronization needed between the parallel tasks must be done via barrier calls.

The primary segments of the execution system operate in a host/slave mode. The master code acts as the host and is initiated first. This code sets up the COMM Network, based on configuration files. Specifically, it first runs a procedure to start a communication server on each computer. The master code next starts all primary segments on the appropriate computers

via the communication library (COMMLIB). The master then executes tasks via COMMLIB within each of the primary segment servers as necessary to complete the desired work for a run. Finally, the segment servers and the COMM Network are shut down.

A graphical user interface (GUI) has been developed to display the state of the FIDO system at all times from start-up to completion of a run. The GUI shows a simplified problem diagram and indicates which parts are starting up, active, inactive, or shutting down by color code.

Data Management

The purpose of the FIDO data manager (DATA in figure 4) is to provide a centralized access service for the storage and retrieval of data during a run of the FIDO system. For a given design problem, the definition of the data to be handled during a run is prepared by the problem designer during the setup phase, which results in several configuration files that define the data in a standardized format and contain initial values for appropriate data. On start-up, the data manager reads and internally stores the information in these files and then enters into a service mode; in this mode, data values are stored and retrieved upon request via COMMLIB calls during a run, and data files are moved among FIDO work and archival directories at the appropriate times. A detailed report and a summary report file are generated to document the data-management activities during a run.

One of the files generated in the setup phase, which contains information about the atomic data elements, provides the basic description of the managed data. In this file, numeric codes are assigned to individual integer, floating point, and character string data elements or to arrays of these. An include file contains corresponding mnemonic macros for each of the numeric codes. These mnemonic macros are used in the discipline driver codes to reference the corresponding data values. Additional information provided for each atomic data element includes the name, type, units, a short description, usage rules, and array length, where appropriate.

Another file identifies the groups of atomic data elements that comprise the actual data packages sent and received by the data manager. This grouping of character strings and integer and floating point numbers results in more efficient use of the message-passing system.

Discipline Segments

The discipline segments are represented by the blocks labeled AERO, PERF, OPTI, INTR, PROP, and STRU around the lower part of figure 4. Each functionally consists of two parts. One part is the set of codes that perform the discipline computations; the other is a driver that handles the discipline communications via COMMLIB and calls on the discipline codes as needed.

The discipline codes generally are established Fortran codes with well-known characteristics and proven reliability. These are the codes that do the actual analysis, which may be computationally intensive. A minimum of changes are made within these codes to prepare them for use with the FIDO system. These changes are those necessary to put these codes into what can be considered subroutine library form. That is, these codes are changed

to make them callable as subroutines with computation and input/output control managed through subroutine arguments.

This control of output from the discipline codes is necessary to the orderly function of the FIDO system. Without control, intermixed output from several disciplines generally scrolls up and off the main workstation window; needed information is lost, overwhelmed by the rarely used detail that most Fortran codes print out. By applying appropriate control, a summary log of the multiple analyses can be kept and the data that are needed are correctly passed from one discipline to another. Data to be kept (including complete output files) are managed by the FIDO system, and the SPY segment is able to retrieve data interactively for display as requested during a run.

The discipline drivers generally are written in C; the discipline codes are called as subroutines. Each driver contains several blocks of code that are invoked by the master as needed to handle the start, analysis, gradient, and exit phases of the problem. Each block contains sub-blocks of code for computation, normal completion and error-handling tasks. The discipline drivers in the FIDO model problem can serve as templates or examples for more complex problems.

The design variable values, file names, and other data are passed through the discipline drivers as arguments to the discipline codes. According to the problem requirements, more than one discipline code may need to be called by a driver. An example is the aerodynamics discipline, which calls separate codes to obtain the induced, wave, and friction drag components in computing total drag for the model problem. The interfaces between the discipline codes and their drivers must be accurately specified in order to provide proper communications. In addition, the user needs to specify what the output files from the code are and how they are to be used by SPY.

For the FIDO model problem, the principal disciplines are aerodynamics, structures, propulsion, and performance. Functionally, the optimizer and interdisciplinary codes are also treated as discipline codes within the system. Because the emphasis in FIDO is on developing the system, the discipline codes were chosen for speed rather than accuracy. They include linear aerodynamics codes and an equivalent plate structures code; these are considered "low fidelity" codes but run in a few minutes on a workstation. In the future they will be replaced with medium- or high-fidelity codes (Euler or Navier-Stokes aerodynamics, finite-element structures), which will take minutes or hours to run on parallel or vector supercomputers.

FIDO Communication Library

The FIDO communication library (COMMLIB in the middle of figure 4) is designed to be used as the communication backbone of a system of codes executed in a heterogenous, distributed network of computers. The various computational tasks (and codes to do these tasks) must be sorted into groups (called segments) for COMMLIB. Each of these tasks is targeted for a different computer. The segments are further sorted into those that are primary and those that are auxiliary. The primary segments are those that perform functions essential

to a run; each requires a driver or server to be started at the beginning of the run. The auxiliary segments are those that do not perform essential functions and so can connect with the COMM Network at any time via a user request. Even for the primary segments, not all the executable code must be loaded with the server; it can be executed later via a forked or remote process if that is more efficient.

The collective execution of the primary segments on a set of computers linked together by a physical network forms a distributed execution system. The part of that system that consists of the communication library, the selected computers, and the physical network is referred to as the COMM Network. The segments use the COMM Network for various communication chores by making calls to COMMLIB.

Although any two segments can directly transfer messages via COMMLIB, a more conservative approach is to use a single data manager. The library includes a special set of message-passing functions that communicate with a data manager segment, which should run in parallel in a service mode while the other segments execute tasks. Thus, the system can be programmed so that nearly all global data (or data used by multiple segments) are sent to and retrieved from the data manager. Both direct messages and data-manager messages can be used as the system designer determines best. Additional special calls are available to send profile information to the profile analyzer segment (currently the data manager) and to send flags to a monitor (currently the SPY segment).

The communication library is merely a tailored interface to a low-level message-passing library (currently the PVM system from Oak Ridge National Labs[1]). It is tailored to provide the necessary functionality in a simple interface format that supports a distributed execution system like that described above. Each call in COMMLIB can actually send several messages to conduct a complex handshaking communication protocol between two segments, which allows the user-level calls to be simple. In addition, this approach provides improved portability and generality to the execution system because only COMMLIB itself needs to be modified if the low-level message-passing software is changed.

SPY Segments

The FIDO SPY segment (SPY Graphics in figure 4) is designed to provide secondary services to system codes that use COMMLIB. These services allow users of the FIDO system to retrieve results while the computations are in progress and display them as text or graphics. In addition, the principal user can alter the values of selected variables in order to provide some guidance to the design process.

Four types of data are available: problem definition, cycle status, cycle history, and profile data. The problem definition data consist of the fixed problem parameters, initial values of some variables, and miscellaneous descriptive information. The cycle status data contain the current cycle number, phase and task; a list of data that is available in DATA and SPY; and miscellaneous timing information. The cycle history data consist of selected scalar values and

[1] PVM information is available by sending the e-mail message "send index from pvm3" to *netlib@ornl.gov*

selected data arrays (on the computational grids) for each completed cycle. Finally, the profile data consist of the execution time histories of the various computational tasks. These data can be displayed on the screen as text or in graphical format. Current graphical formats include line plots (e.g., cycle history of the objective function) and contour plots of distributed data (e.g., surface pressures (as in figure 5), stresses, or deflections).

Multiple instances of the SPY code can connect to and communicate with the COMM Network that is used by a set of primary segments. Thus, the master/slave paradigm is augmented to include multiple (but limited) masters. The SPY connections can be initiated at any time while the FIDO system is running a problem. In the current implementation, new SPY's become active only at the beginning of a new design cycle; future implementations will allow activity to begin within a cycle.

Actually, two types of SPY segment exist. One is designated as the 'Master SPY' and has special properties. It can be invoked only by the user who is running the FIDO problem (the "designer"); only this user is allowed to change the current values of design variables, constraints, and initially-set parameters. This capability provides some measure of manual control over the design process. Because the designer may need advice during the design process (perhaps from a remotely located expert in one of the disciplines involved), the other type of SPY segment is provided.

This SPY type is designated as an 'Agent SPY'; it is limited to displaying data, but can exist in multiple instances. It can be invoked with proper permission from any location connected to the network (e.g., anywhere on the Internet). Thus, the remote expert's workstation can display any of the information available to the designer (shown schematically in figure 6). In fact, several users can examine the same or different FIDO output at different locations at the same time.

CONCLUDING REMARKS

The Framework for Interdisciplinary Design Optimization (FIDO) is being developed as a general programming environment for automating the distribution of complex computing tasks over a networked system of heterogeneous computers. The FIDO system facilitates communications between computational tasks distributed over a computer network and provides for automatic interactions in multidisciplinary problems, for example, how to reach a nearly optimal consensus in the aircraft design process.

In the FIDO system, the computers involved are networked together, have access to centralized data, and work on their parts of the design simultaneously in parallel whenever possible, under the direction of a master code. Each computational task can be assigned the computer type most appropriate for it. An auxiliary code, SPY, is provided for viewing results as they are produced and for steering the design process. In the viewing mode, it can be run as multiple instances and from remote locations. The FIDO software is written in a modular form in order to ease migration to upgraded or completely new problems: different codes can be substituted for each of the current code modules with little or no effect on the others.

The FIDO system has been designed to be adaptable to any distributed computing problem. It has been demonstrated for a simplified High-Speed Civil Transport (HSCT) aircraft design problem; currently, a much more complex HSCT problem is being implemented. The FIDO system's potential for commercial use rests in the capability it provides for automatically coordinating diverse computations on a networked system of workstations and computers. For example, FIDO could provide the coordination required for the design of fixed- or rotary-wing aircraft, automobiles, ships, spacecraft, computer systems, or electronics. It could also be used for modeling large-scale interacting systems, such as in economic or ecological models.

APPENDIX. FIDO PROJECT SUMMARY

FIDO System Features:

Provides synchronous control in a heterogeneous computing environment
Is adaptable to a wide variety of problems that use multiple computer programs
Has modular form to ease migration to new problems
Incorporates a variety of debugging features
Includes SPY capability for monitoring computational progress, displaying variables graphically, and modifying variables for design steering

FIDO Project Accomplishments:

Optimization of simple HSCT design problem
Use on network of UNIX® workstations (Sun, SGI, DEC)
Demonstration of monitoring, graphics display, and variable modification through SPY
Development and demonstration of flexible file-based data management

Future FIDO Plans:

Demonstrate FIDO on a complex HSCT design problem (NASA Langley HiSAIR "Pathfinder")
Replace low-fidelity-analysis with high-fidelity-analysis computer programs
Improve graphical user interfaces for problem setup and for control
Fully document the software and publish findings
Get and apply feedback from Beta Tests

FIDO Point of Contact:

Dr. Robert P. Weston
Computational Sciences Branch
Mail Stop 159
NASA Langley Research Center
Hampton VA 23681–0001

(804) 864–2149
FAX: (804) 864–6134
r.p.weston@larc.nasa.gov

FIDO Team:

FIDO Concept Originator and Project Designer:

— Dr. Thomas M. Eidson, High Technology Corporation

FIDO Project Leader:

— Dr. Robert P. Weston, Computational Sciences Branch

Current FIDO Working Group:

— Ramki Krishnan, Analytical Services & Materials, Inc.
— Jim Townsend, Computational Sciences Branch
— Ben James, Lockheed Engineering & Sciences Co.
— Kelvin Edwards, Analytical Services & Materials, Inc.
— Mary Adams, Computational Sciences Branch
— Ray Gates, Computer Sciences Corp.
— Don Randall, Computer Systems Branch

Other FIDO Contributors:

— Peter Coen, Vehicle Integration Branch
— Tom Zang, Computational Sciences Branch
— Larry Green, Computational Sciences Branch
— Gary Giles, Aeroelastic Analysis & Optimization Branch
— Gus Dovi, Lockheed Engineering & Sciences Co.
— Greg Wrenn, Lockheed Engineering & Sciences Co.
— Kim Cunningham, Computer Sciences Corp.
— Bill LaMarsh, Unisys Corp.
— Nina Hathaway, Unisys Corp.
— Tom Crockett, ICASE
— Bill von Ofenheim, Flight Software & Graphics Branch

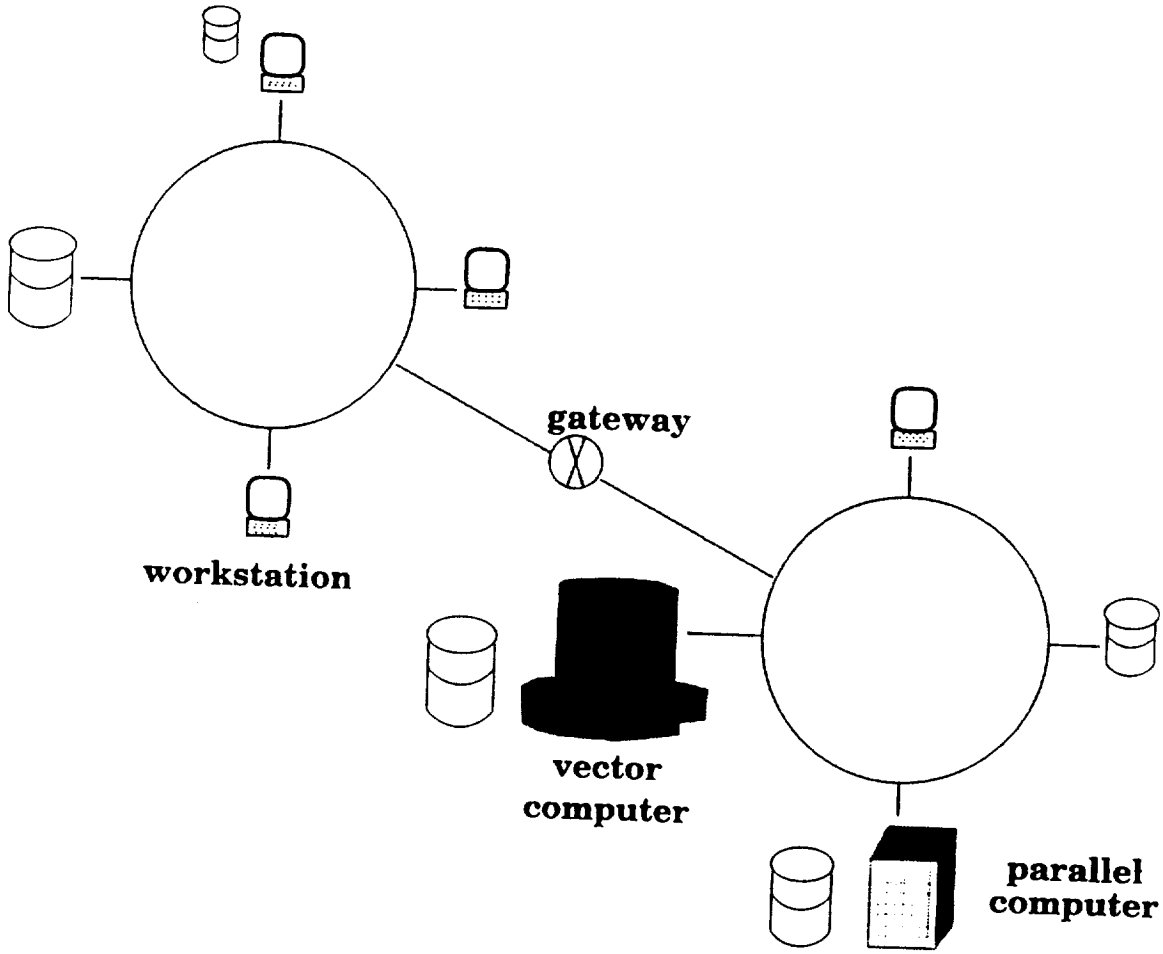

Figure1.– Heterogeneous computing environment.

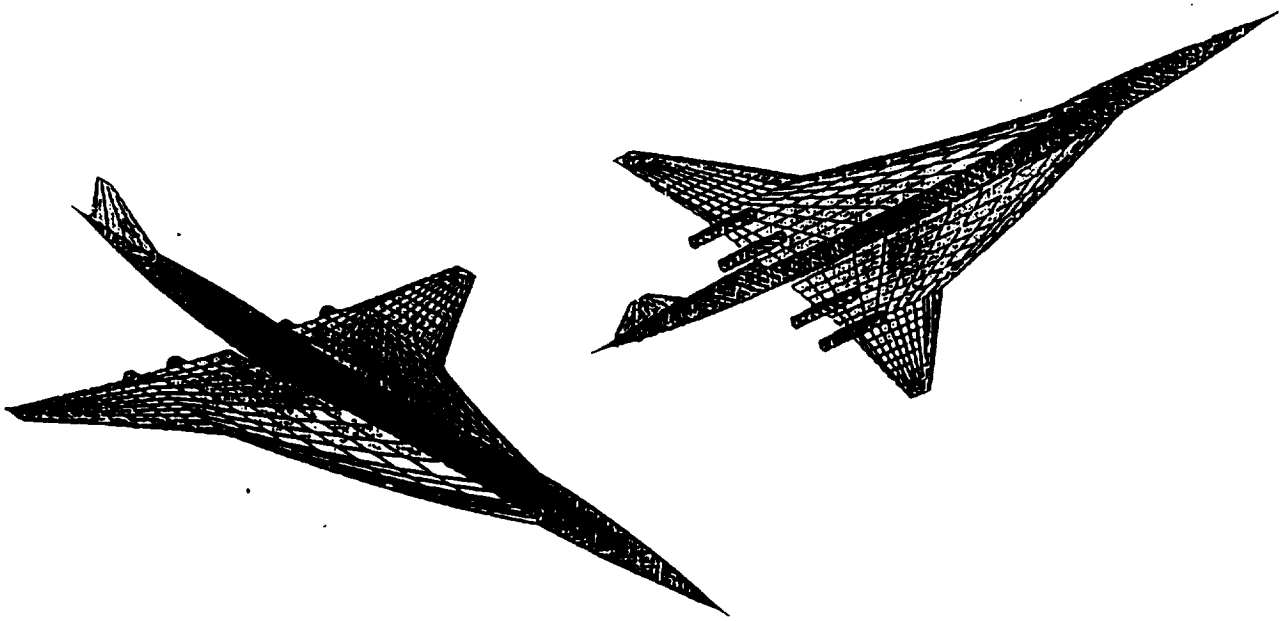

Figure2.– NASA Langley Research Center HSCT baseline configuration.

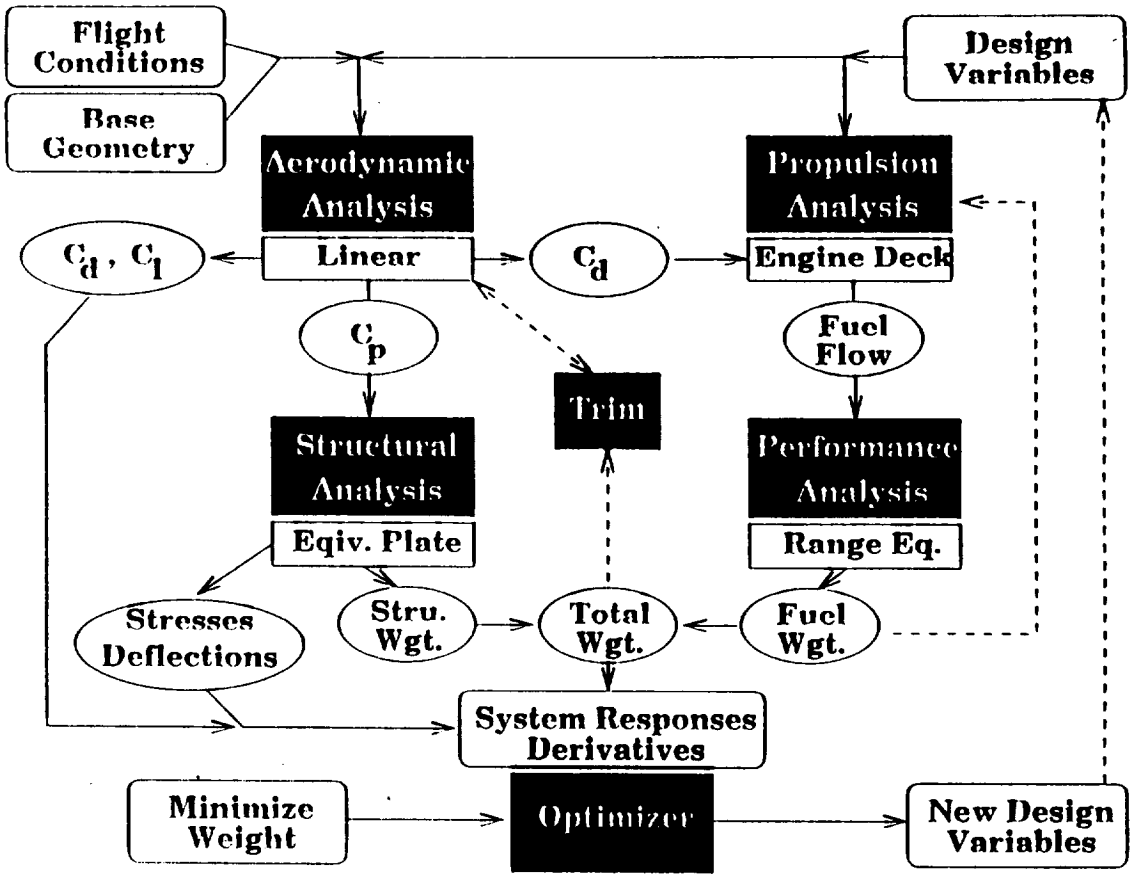

Figure 3.– FIDO model diagram for HSCT design problem.

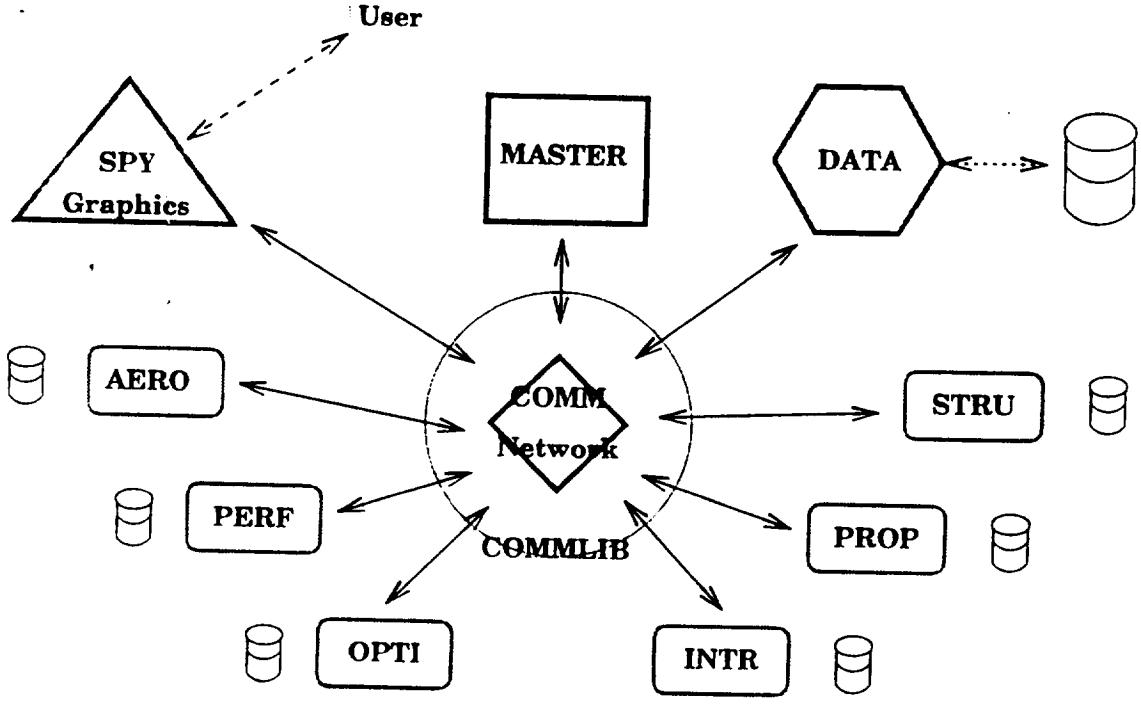

Figure 4.– Schematic of FIDO system interactions.

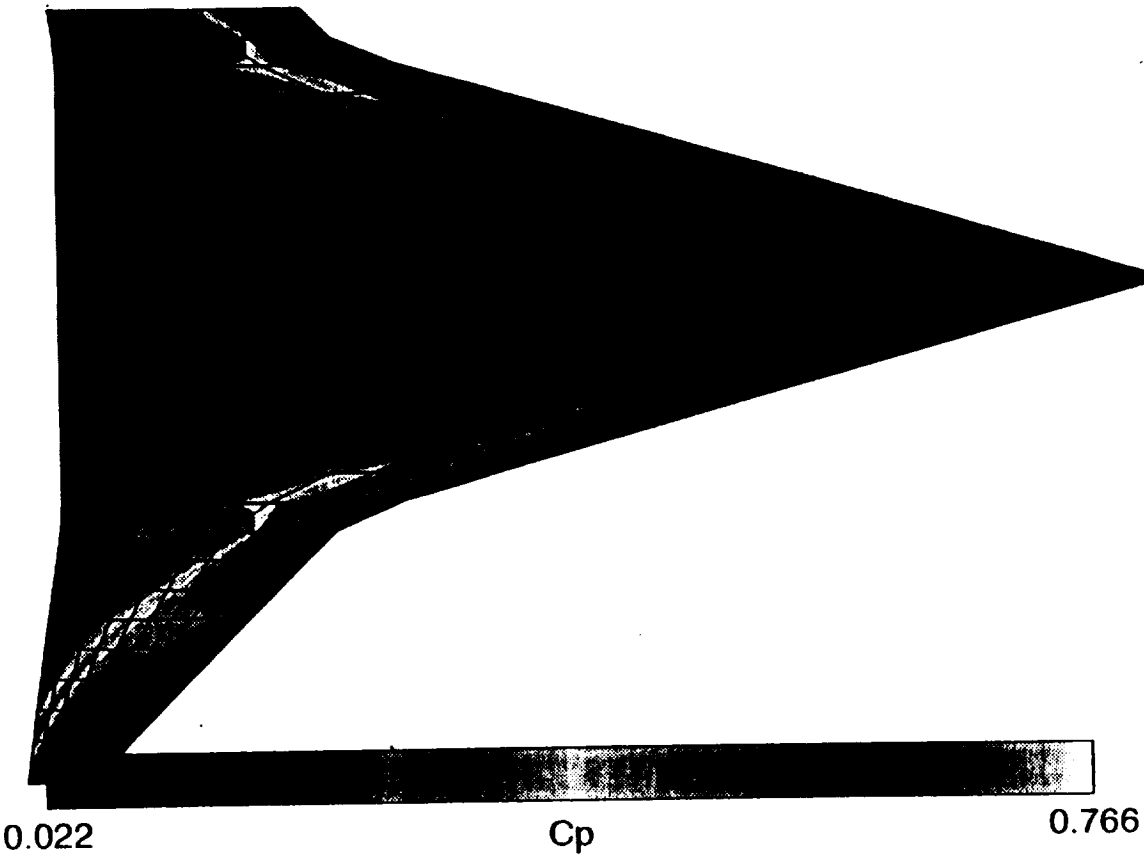

Figure 5.– Sample HSCT surface contour plot (original in color).

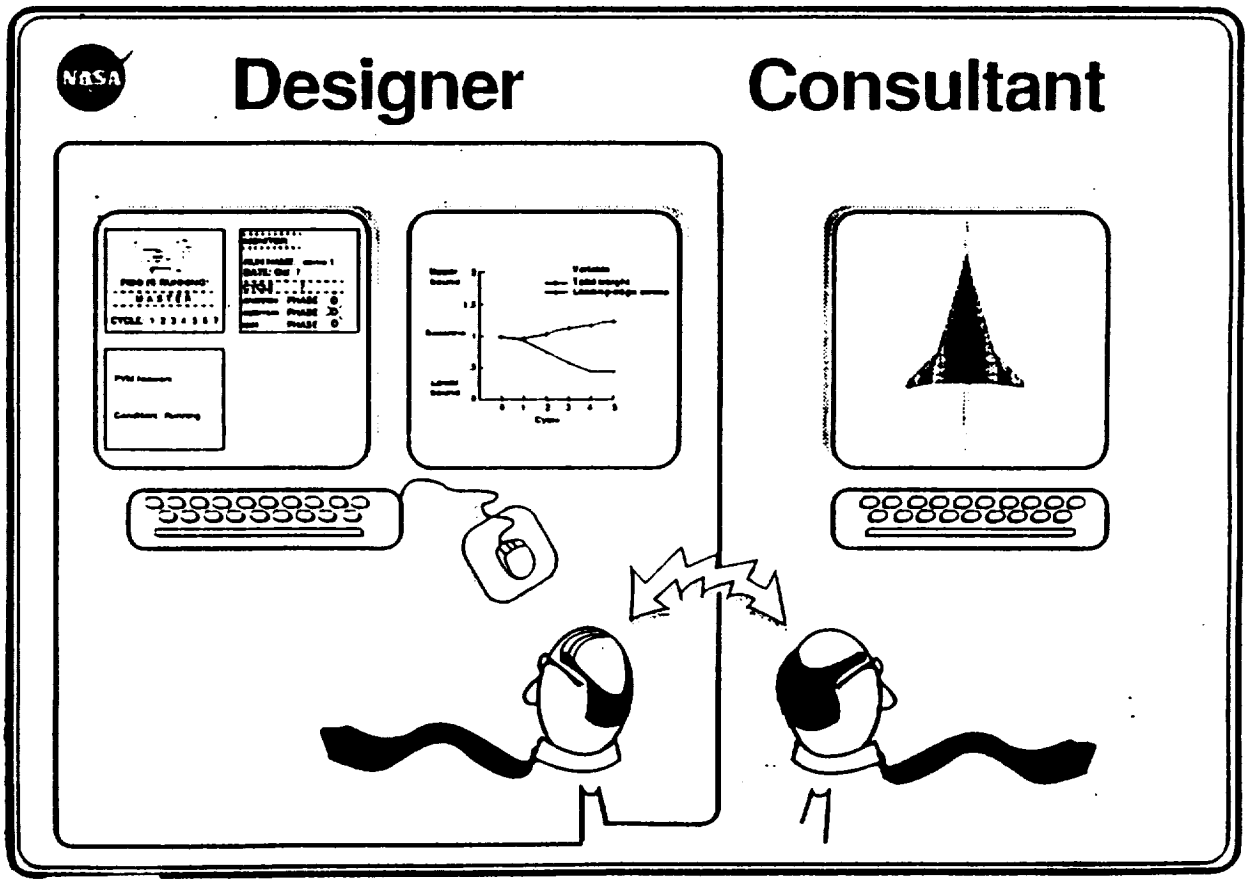

Figure 6.– Concept of designer interaction with consultant through SPY.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	December 1993	Technical Memorandum

4. TITLE AND SUBTITLE A Programming Environment for Distributed Complex Computing. An Overview of the Framework for Interdisciplinary Design Optimization (FIDO) Project

5. FUNDING NUMBERS
505-59-53-08

6. AUTHOR(S) James C. Townsend, Robert P. Weston, Thomas M. Eidson

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
NASA Langley Research Center
Hampton, VA 23681-0001

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Washington, DC 20546

10. SPONSORING / MONITORING AGENCY REPORT NUMBER
NASA TM-109058

11. SUPPLEMENTARY NOTES

Townsend and Weston: Langley Research Center, Hampton, VA; Eidson: High Technology Corporation, Hampton, VA.
NASA Langley Technology Opportunities Showcase (TOPS) Exhibit H120b.

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Unclassified - Unlimited

Subject Category 61

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The Framework for Interdisciplinary Design Optimization (FIDO) is a general programming environment for automating the distribution of complex computing tasks over a networked system of heterogeneous computers. For example, instead of manually passing a complex design problem between its diverse specialty disciplines, the FIDO system provides for automatic interactions between the discipline tasks and facilitates their communications.

The FIDO system networks all the computers involved into a distributed heterogeneous computing system, so they have access to centralized data and can work on their parts of the total computation simultaneously in parallel whenever possible. Thus, each computational task can done by the most appropriate computer. Results can be viewed as they are produced and variables changed manually for steering the process. The software is modular in order to ease migration to new problems: different codes can be substituted for each of the current code modules with little or no effect on the others.

The potential for commercial use of FIDO rests in the capability it provides for automatically coordinating diverse computations on a networked system of workstations and computers. For example, FIDO could provide the coordination required for the design of vehicles or electronics or for modeling complex systems.

14. SUBJECT TERMS			15. NUMBER OF PAGES
Heterogeneous computing Distributed computing Multidisciplinary design optimization Design automization	Communications library		13 16. PRICE CODE A03
17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
unclassified	unclassified	unclassified	